Unidirectional Quorum-based Cycle Planning for Efficient Resource Utilization and Fault-Tolerance

The final publication is available at IEEE via http://dx.doi.org/10.1109/ICCCN.2016.7568595

Cory J. Kleinheksel, Member, IEEE and Arun K. Somani, Fellow, IEEE
Electrical and Computer Engineering, Iowa State University, Ames, Iowa 50011
Telephone: (515) 294–0442, Fax: (515) 294–9273, Email: {cklein, arun}@iastate.edu

arXiv:1609.07781v1 [cs.NI] 25 Sep 2016

***Abstract*—In this paper, we propose a greedy cycle direction heuristic to improve the generalized R redundancy quorum cycle technique. When applied using only single cycles rather than the standard paired cycles, the generalized R redundancy technique has been shown to almost halve the necessary light-trail resources in the network. Our greedy heuristic improves this cycle-based routing technique's fault-tolerance and dependability.**

For efficiency and distributed control, it is common in distributed systems and algorithms to group nodes into intersecting sets referred to as quorum sets. Optimal communication quorum sets forming optical cycles based on light-trails have been shown to flexibly and efficiently route both point-to-point and multipoint-to-multipoint traffic requests. Commonly cycle routing techniques will use pairs of cycles to achieve both routing and fault-tolerance, which uses substantial resources and creates the potential for underutilization. Instead, we use a single cycle and intentionally utilize R redundancy within the quorum cycles such that every point-to-point communication pairs occur in at least R cycles. Without the paired cycles the direction of the quorum cycles becomes critical to the fault tolerance performance. For this we developed a greedy cycle direction heuristic and our single fault network simulations show a reduction of missing pairs by greater than 30%, which translates to significant improvements in fault coverage.

I. Introduction

Internet backbones and data centers have become key pieces of national infrastructure rising almost to the level of roads, bridges, electricity, and water. Optical networks in these settings are depended upon for high speed communications. Failures within a network are to be expected and can happen as much as every couple days [1].

Knowing the unicast or multicast requests a priori is often not possible. This constraint makes protection against faults in those arbitrary communication paths a challenge. An efficient all node pairs protection scheme supporting both unicast and multicast communication is necessary.

For efficiency and distributed control, it is common in distributed systems and algorithms to group nodes into intersecting sets referred to as quorum sets. Quorum-based cycle routing can efficiently support arbitrary point-to-point and multi-point optical communication [2]. Cycles are created using quorums of nodes. Within a cycle, multicasts to all nodes in that cycle are possible. The proof that cyclic quorums have an "all-pairs" property guarantees unicast capabilities in a network using the quorum cycles routing technique [3]. Efficient broadcasts can be achieved with $O(\sqrt{N})$ multicasts by exploiting the same properties.

Fault tolerance of optical networks with quorum-based cycle routing has been analyzed [4]. Observing communication pairs were being protected by more than one cycle, R redundancy quorums were used to improve resource efficiency [5]. Resource efficiency can come at a cost and in this research, we propose a heuristic to improve the redundant quorum cycle routing techniques. Our single fault network simulations show a reduction of missing pairs by 31.16 - 48.74% and 35.48 - 44.85%, $R = 2$ and $R = 3$ respectively, which translates to significant improvements in fault coverage.

The next sections describe the network model, light-trails, and cycle routing. Section IV establishes how quorums can be used for optical network routing. We analyze the current strengths and opportunities for improvements to R redundant quorum-based single cycle routing in Section V. Lastly in Section VI, we present a greedy cycle direction heuristic to improve the optical network fault tolerance.

II. Network Model

The fiber-optic networks consist of several transmitters and receivers (nodes) interconnected by fiber-optic cables. The cables form the links (i.e., edges) between those nodes, which leads to a convenient model of a network in terms of a graph $G = (V, E)$. V are the set of nodes in the network and E are the set of edges.

Edge (e_i, e_j) is a fiber-optic link connecting nodes e_i and e_j in the network, where $e_i, e_j \in V$ and $(e_i, e_j) \in E$. It is a general assumption that the same set of optical wavelengths are available on all edges in E. The number of wavelengths available per optical fiber is dependent on the fiber-optic cables and the transmitter/receiver pairs.

A. Light-Trails

Lightpaths were a critical building block in the first optical communications, but required significant traffic engineering and aggregation to support point-to-point communication, or pay the penalty of low resource utilization on the fiber-optic link. Light-trails were proposed [6] as a solution to the challenges facing lightpaths and could be built using commercial off-the-shelf technology. Significant contributions have been made to enable adoption and advance the light-trail architecture [7]–[9]. Light-trail communication is all optical and uses the same wavelength(s) from start to end node;

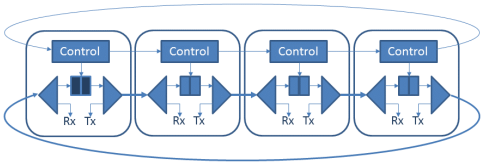


Figure 1. Cycle formed using the light-trail architecture

avoiding Optical-to-Electrical-to-Optical (O/E/O) conversions, energy inefficiencies and time delays at intermediate hops.

Light-trails enable fast, dynamic creation of a unidirectional optical sharable communication channel. This communication channel allows for channel receive and transmit access to all connected nodes, making them more suitable for IP-centric traffic [7]. Communications from an upstream node to one or more downstream nodes can be scheduled. The scheduling protocol avoids collisions within a light-trail and controls when nodes are able to transmit to downstream nodes. The scheduling is generally assumed to occur over a control channel, which may or may not be separate from the shared optical fiber that is being used for the light-trail.

Optical shutters allow for wavelength reuse within the network. Start and end nodes have their optical shutters in the *off* state, while intermediate nodes have their optical shutters in the *on* state. This effectively isolates an optical signal to a specific light-trail and allows for reuse of optical wavelength(s) elsewhere in the network.

Nodes can receive from the incoming signal while the signal is simultaneously continuing to downstream nodes, sometimes referred to as a drop and continue function. Early technology supported only a few wavelengths; however, the latest devices may support over 100 channels, hence allowing multiple light-trails to share the same edge in the network for a combined over 1-Terabits/s [10].

III. Light-Trails, Cycle Routing, and Fault Tolerance

Point-to-point and multi-point traffic requests have a set of nodes $C = \{e_i, ..., e_j\}$ that wish to communicate and need to be protected against network faults. Establishing a primary and backup multicast path from every node to every other node in C can be a waste of resources. We utilize the light-trail architecture in the form of a cycle (Fig. 1). The cycle will both route the multi-point request and protect it at the same time using fewer resources.

Figure 1 is a light-trail where the start and end node is the same node, referred to as the *hub node*. The hub node has its optical shutters in the *off* state, while intermediate nodes have their optical shutters in the *on* state.

Failures within an optical network are to be expected. SONET rings can be used to protect point-to-point and shared paths while enabling failure location. The generalization of p-cycle protection to allow for path and link protection was proposed in [11]. P-cycle protection of unicast and multicast traffic networks requires preconfiguration, and the offline nature allows for the efficient cycles to be selected. The Optimized Collapsed Rings (OCR) single link protection heuristic was developed to address the heterogeneous, part multicast / part unicast, nature of WDM traffic [12].

The multi-point cycle routing algorithm (MCRA) uses bidirectional cycles for fault tolerance and is capable of supporting SONET rings and p-cycles [1]. ECBRA is a significant improvement of MCRA and outperforms the OCR heuristic [9]. ECBRA heuristic balances optimality and speed, taking $O\left(|E|\,|C|^3\right)$ steps to find a close to optimal cycle.

IV. Quorums Sets for Routing

Quorums have been used for efficient point-to-point, multicast, and all-to-all traffic requests in optical networks [2]. This is important because traffic in many optical networks is heterogeneous meaning the routing framework must be able to handle all types.

Quorums sets cover N entities, in this case $N = |V|$ optical network nodes. The number of entities, N, also defines the number of small subsets, i.e., quorums, that will be used in our solution.

A quorum set minimally has the property that all quorums must intersect. Specifically for distributed implementations, it is also desirable that each quorum have equal work and equal responsibility within the quorum set. Cyclic quorum sets have these properties [13] and we use them for efficient communication.

Point-to-point, multicast, and all-to-all traffic can be routed through an optical network with N cycles based on cyclic quorums [2], [4]. Cyclic quorums sets are used as the basis for cycle routing to guarantee all possible node pairs occurred in at least one cycle [3]. When performing a multicast, if participants belong to the same cycle, then one cycle can be used. Realistically, requests will span multiple quorums and/or be larger than a single quorum cycle. The worst case becomes a broadcast with an upper bound of requiring no more than k cycles to route and protect broadcast traffic.

A. Defining Cyclic Quorum Sets

Cyclic quorum sets are based on cyclic block design and cyclic difference sets. However, searching for optimal sets requires an exhaustive search [13]. Cyclic quorum sets are unique in that once the first quorum (Eq. 1) is defined the remaining quorums in the set can be generated via incrementing the indices (modulus to keep indices within bounds is not shown in Equation 2 for conciseness).

$$S_0 = \{e_0, \ldots, e_j\} \tag{1}$$

$$S_i = \{e_{0+i}, \ldots, e_{j+i}\}\,,\ \forall\, i \in 0, 1, \ldots, N-1 \tag{2}$$

For simplicity, assume $e_0 \in S_0$ without loss of generality (any one-to-one re-mapping of indices can result in this assumption).

For our work, we used the $N = 4, \ldots, 111$ optimal cyclic quorums from [13].

The final publication is available at IEEE via http://dx.doi.org/10.1109/ICCCN.2016.7568595

B. Redundant Cyclic Quorums Sets

In this section we define R redundant cyclic quorums sets as proposed in [5]. Quorum-based cycle routing solutions used cyclic quorums to form a set of communication cycles which were shown to be almost fault tolerant in fiber optic networks [2], [4].

There are N quorums in a quorums set and each quorum has k nodes. In analysis of networking capabilities, we are interested in whether every node can communicate with every other node (all-pairs). Equation 3 considers the number of pairs within a quorum, e.g., the pairs made between k nodes communicating with $(k-1)$ other nodes in a single quorum. Equation 4 considers the total pairs formed by all N quorums in the set. For convenience we set M to be the total pairs for a given network with N nodes and k optimal quorum size.

$$\frac{k(k-1)}{2} = O(k^2) \tag{3}$$

$$N\frac{k(k-1)}{2} = O(Nk^2) \tag{4}$$

$$M = O(Nk^2) \tag{5}$$

When the quorum size, k, is minimal or larger, every pair of nodes (e_i, e_j) occurs together within a quorum in the set at least once. Optical networking, however, requires all directional point-to-point pairs to exist, i.e., both pairs (e_i, e_j) and (e_j, e_i). Previously this had been addressed by pairing each cycle with the same cycle and its direction reversed.

It was observed in [2] that the quorum-based cycle routing solution had some node pairings occurring together in multiple cycles and it was proposed that these could be used for load balancing. As an alternative to that option, [5] added a requirement that every pair (e_i, e_j) would occur together within at least R quorums rather than just one. Exploiting the natural occurrence of redundant pairs is an attempt to eliminate the need for paired cycles, thus moving the redundancy from the paired cycles and putting the redundancy in the quorums.

The number of quorums in the solution remained the same N, hence to create the additional pairs the quorum size had to be enlarged to $\hat{k}$. Equation 6 calculates the number of node pairs in quorums of size $\hat{k}$. Equation 7 is our requirement that the total number of pairs have increased R times from the original total pairs, M. Finally, Equation 8 solves for size $\hat{k}$ in relation to optimal k.

$$N\frac{\hat{k}(\hat{k}-1)}{2} = O(N\hat{k}^2) \tag{6}$$

$$O(N\hat{k}^2) = RM \tag{7}$$

$$\hat{k} \approx \sqrt{R}k \tag{8}$$

This result is powerful. The number of node pairs increased by R times, but the size of k only increased by a factor of $\sqrt{R}$. Using this reduced growth rate to our advantage, many node pairs can be created without substantially increasing the resources used. Additionally, this growth rate is far slower than simply duplicating a cycle, hence significantly challenging the need for paired cycles and opening the door for considerable resource savings.

To the best of our knowledge, no efficient algorithm is known to find quorums of minimum size, particularly with the additional requirement that entity pairs appear a minimum R times within the quorums set solution. The authors in [13] used a brute force search to find optimal cyclic quorums for $N = 4 \ldots 111$. Using our generalized result from Eq. 8, we too used a brute force search beginning with the smallest possible quorum size for a given number of nodes N and a given desired redundancy factor R.

The resulting redundant quorums were utilized in following section as we analyzed the efficiency and fault tolerance of quorum-based cycle routing in optical networking.

V. Redundant Cyclic Quorums Set - Single Cycle Network Analysis

Optical networks are highly depended upon. The fault tolerance aspect of these route designs are important. In [2], [4], ECBRA was used to route each of the quorums-based cycles. It was shown that the quorums set approach provided fault tolerance and [2] showed that this technique required far fewer links to accomplish the routing of all-to-all traffic, when compared to using point-to-point connections.

As an unintended benefit, some quorums sets resulted in node pairs occurring in more than one quorums-based cycle. It was these occurrences of node pairs multiple times that improved the fault tolerance performance.

Redundant node pairs can also be generated intentionally as described in Section IV-B. These can be used to eliminate the paired cycle implementations and use only a single cycle, significantly reducing the required amount of network resources while still maintaining a similar level of fault-tolerance. Our work reported in [5] examined this approach.

In this section, we analyze the strengths to deliver resource efficient routing solutions, while identifying opportunities to improve fault tolerance in the network.

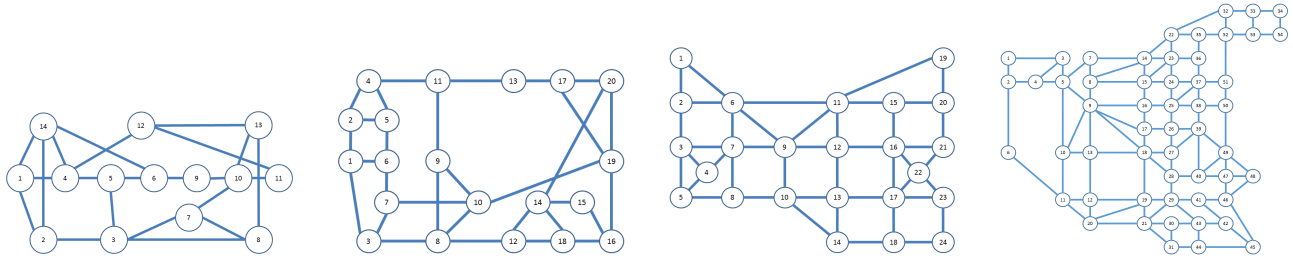
A. Fault Model

The fault model assumed for our work is the link (edge) failure. While a simple model, it does cover most real single fault scenarios.

The most direct fault to consider is the optical link fault. This occurs when a link is broken, like planned maintenance or the accidental severing during land excavation. Modeling link faults as a single edge failure is straightforward.

Each modeled node needs a pair of transmitters and receivers for each occurrence in a cycle. These pairs of devices can fail too. Short of a natural disaster, pairs will likely fail independently of one another. When a transmitter/receiver pair fails within a modeled node, the affect on the global network is similar to that link failing. Modeling as a single edge failure, while not an exact fault mapping, is an appropriate abstraction.

The final publication is available at IEEE via http://dx.doi.org/10.1109/ICCCN.2016.7568595



(a) (b) (c) (d)

Figure 2. Networks used for simulations: Figure (a) NSFNET, 14-Node/22-Link, Figure (b) ARPANET, 20-Node/31-Link, Figure (c) American Backbone [14], 24-Node/43-Link, and Figure (d) Chinese Backbone [14], 54-Node/103-Link.

Table I
MEAN LINKS USED BY SINGLE CYCLES COMPARED TO PAIRED CYCLES USING REDUNDANT CYCLIC QUORUMS (95% CI)

	$\mathbf{R = 1}$ **(Paired)**	$\mathbf{R = 2}$ **(Single)**		$\mathbf{R = 3}$ **(Single)**	
Network	**Links**	**Links**	**Reduction (%)**	**Links**	**Reduction (%)**
NSFNET	249.32 ± 1.37	135.41 ± 0.60	-45.69	145.05 ± 0.74	-41.82
ARPANET	511.90 ± 1.87	269.63 ± 0.76	-47.33	294.75 ± 0.96	-42.42
American	641.38 ± 2.10	360.03 ± 0.98	-43.87	376.86 ± 0.77	-41.24
Chinese	2673.30 ± 7.11	1527.32 ± 3.50	-42.87	1635.70 ± 2.99	-38.81

B. Experimental Setup

Maintaining the ability to serve all dynamic point-to-point traffic requests despite fault is important. We examined the fault tolerance of the NSFNET, ARPANET, American backbone, Chinese backbone networks (Fig. 2).

To model the fault, we simulate the failure of each edge, $(e_i, e_j) \in E$, in the network model, $G = (V, E)$. We then examine the network's ability to serve all potential point-to-point requests by counting pairs of nodes that would be able to communicate and conversely those pairs that are unable to communicate.

C. Network Routing Efficiency and Fault Tolerance

We examine redundant quorums by using additional cyclic quorum redundancy with just a single cycle compared to paired cycles used in prior art. The additional redundancy is used to distribute node communication pairs across different cycles, e.g., $0 \rightarrow 1$ and $0 \leftarrow 1$ may not occur within the same cycle. Recall that optical light-trail cycles are unidirectional, so prior art would satisfied both $0 \rightarrow 1$ and $0 \leftarrow 1$ with a pair of cycles. One cycle would form the forward communication $0 \rightarrow 1$, while the second cycle would form the backward communication $0 \leftarrow 1$. By distributing the communication pair responsibility throughout the network, we may only need a single cycle, thus significantly reducing the required amount of network resources.

We used four common networks (Fig. 2) and an implementation of the ECBRA heuristic [9] to perform the cycle routing. ECBRA is sensitive to node and edge numbering that a total of 100 variations on the inputs were considered, each being a one-to-one mapping with the respective network. For simulation of prior art, we used the $N = 4, \ldots, 111$ optimal cyclic quorums from [13]. Redundant cyclic quorums for $R = 2$ and $R = 3$ were found using the techniques described in Section IV-B and in our work reported in [5].

1) Fault-free operational analysis: It is expected that a majority of the time the optical network will be operating without faults. It is important that the resource utilization during this period be analyzed.

The metric we use to measure resource utilization is the number of links used in a solution. Comparing network-to-network is not particularly insightful, but comparing multiple solutions for a particular network is. The more links that a set of quorum cycles use, the fewer (wavelength) resources that can be assigned to each link. Additionally, each logical link represents a required physical transmitter and receiver, hence capital costs.

Table I shows significant 38.81 - 42.42% resource reduction when using $R = 3$ redundancy in quorums over the more traditional, prior art methods using paired cycles. $R = 2$ gives even better resource reduction. This reduction represents the potential for lower capital costs in terms of physical transmitters and receivers needed and the potential for more (wavelength) resource availability within the network. The paired cycles results with a 95% confidence interval (CI) for $R = 1$ is in column two of Table I for comparison to the single cycle, increased quorum redundancy technique. The technique uses far fewer links (shown in columns 3 - 6.) $R = 2$ comes close to halving the necessary resources.

Previously paired light-trails were used to form all of the point-to-point communication node pairs. We analyze the impact of increasing the R redundancy within quorums and its impact of keeping resource utilization low using the missing node pairs metric.

Ideally, like the paired cycle case, there would be 0% missing, however single cycles do not have the benefit of both

The final publication is available at IEEE via http://dx.doi.org/10.1109/ICCCN.2016.7568595

(e_i, e_j) and (e_j, e_i) pairs occurring in the same cycle. The dramatic reduction in resource utilization came at a trade off of a few missing communication pairs. $R = 2$ missed 0.95% or fewer on average (95% CI), and $R = 3$ missed even fewer at 0.21% or less on average (95% CI). Redundant $R = 3$ cycles performs approximately 2+ times better every time. As seen in Table I, this performance improvement came at a slightly higher cost, while still being significantly smaller than the state of art approach.

It is the limitations of unidirectional optical light-trail cycles with its required one optical shutter in the *off* state per cycle that has caused the missing pairs and the potential need for additional compensation steps. Compensation is possible using an off-the-shelf solution of an additional routing step involving an Optical-to-Electrical-to-Optical (O/E/O) conversion and retransmission by a hub node. Even so, on average the $R = 2$ and $R = 3$ redundant quorums cycle solutions would require infrequent additional steps considering the missing pairs are less than 1% on average.

2) Fault-tolerant operational analysis: Using the generalized R quorum redundancy rather than cycle pairs can save significant resources; however, this can come at a significant determent to fault tolerance.

To model the fault, we simulated the failure of each used edge, $(e_i, e_j) \in E$, in 100 node mappings of each network model, $G = (V, E)$. The edges not used in a particular mapping are ignored to prevent biasing the results with zero missing pairs data points. We then examine the network's ability to serve all potential point-to-point requests by counting pairs of nodes that would be able to communicate and conversely those pairs that are unable to communicate. The results are then reported as fault coverage. To calculate fault coverage in this scenario, we calculated the mean connected pairs divided by the total pairs.

$$1 - \frac{Mean\ Missing\ Pairs}{Total\ Pairs} \tag{9}$$

100% would be perfect coverage, whereas 0% would be no fault coverage at all.

Our simulation results (Fig. 3) shows the redundant quorum-based cycle technique had 96.52 - 99.36% and 97.81 - 99.71% fault coverages, $R = 2$ and $R = 3$ respectively, in the four networks tested. We compare the state-of-art paired cycle approach with our quorum redundant technique with single cycles that uses significantly fewer resources. With single link failures, the paired cycles had a mean missing communication pair rate of less than 3 pairs or less than 0.53% across all networks (95% CI). Hence the $R = 1$ (Paired) column shows mean fault coverage percentages is greater than 99.47% for all four networks. The redundant quorum cycles technique, $R = 2$ and $R = 3$ (Single), could not reach that level of coverage, but did achieve a mean fault coverage rate (95% CI) of greater than 96.52 and 97.81%, respectively, across all networks.

While neither single cycle $R = 2$ or $R = 3$ could achieve the same level of fault coverage as the paired cycle solution, they did have missing pair rates better than 3.48

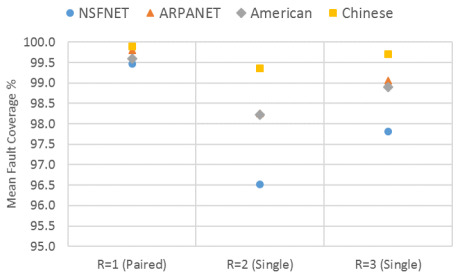


Figure 3. Percent mean fault coverage of the single cycle, redundant quorum solution experiencing a single link fault.

and 2.29%, respectively, while achieving significant resource savings. In networks where an additional approximately 40% of resources could better be utilized for communication rather than redundancy, the trade off of missing a relatively small percentage of communications during fault conditions may be considered tolerable.

VI. Improving Single Cycle Routing Based on Redundant Cyclic Quorums

In Section V-C1, we had an interesting observation that although the cyclic quorums set method was proven to guarantee that all of the node pairs exist, the implementation that utilized the quorums still had the limitation of unidirectional light-trail cycles requiring one optical shutter in the *off* state per cycle. This led to missing pairs even in fault-free states of the network and likely contributed to at least some of missing pairs during fault scenarios.

When we used paired cycles, the routing order within a cycle did not matter. In any given pair of cycles, there was one in the forward direction and the other in the backwards. For example, if the forward cycle route formed the $0 \rightarrow 1$ communication route, then the backward cycle route would form the $0 \leftarrow 1$ communication route.

With a single, unidirectional cycle we were taking for granted that redundant cyclic quorums were guaranteeing two pairs of all node pairs. However, the route direction of the pairs is not enforced. Hence, without the backwards cycle from paired cycle, it turned out that the single cycle could still result in two forward $0 \rightarrow 1$ paths in the network's cycle route implementation rather than one of each ($0 \rightarrow 1$ and $0 \leftarrow 1$ communication routes).

Looking closer at the source of the problem, i.e., the unidirectional optical light-trail cycles, we developed an alternative heuristic solution to improve the performance. The cycles were routed using the ECBRA heuristic [9] and the redundant cyclic quorums for $R = 2$ and $R = 3$. Once routed, we treated the first node in the corresponding cyclic quorum to be the hub node and then formed the cycle in the order provided by the ECBRA heuristic. By controlling the order to either be forward or backwards on a particular unidirectional optical light-trail cycle, we can influence which node pairs are formed.

The final publication is available at IEEE via http://dx.doi.org/10.1109/ICCCN.2016.7568595

A. Greedy cycle direction based on missing pairs

There are $|V| = N$ redundant cyclic quorums for a given network $G = (V, E)$. Each of these quorums has a corresponding cycle route and each route is either in the forward or backward direction. This results in $O\left(2^N\right)$ possible combinations of cycle directions for a network.

Next, we describe our greedy algorithm to determine cycle direction and evaluate the improvement in network performance. Ultimately, by controlling the forwards or backwards direction of a cycle, we want to decrease (and eliminate) the number of missing pairs under fault-free conditions. While doing this, the network's fault tolerance is also anticipated to increase.

Algorithm 1 greedily chooses each cycle's initial direction. All forward and backward node pairs need to be formed in the network. We keep track of how many of each pair have been formed in variable PC on Line 1. We iterate through all of the routed cycles, one for each redundant quorum (Line 2). The cycle's direction is chosen by which direction will eliminate more missing pairs from PC. Then all of that direction's pairs are added to PC (Line 7 or Line 10). The next cycle's direction is chosen the same way until all cycles have been assigned an initial direction.

Algorithm 1 Initial Cycle Direction($Cycles$,V)

```
1:  Pair Count PC[e_i][e_j] ← 0, ∀e_i, e_j ∈ V
2:  for all c ∈ Cycles do
3:      Count number of new pairs added to PC if Cycle c.direction = Forward
4:      Count number of new pairs added to PC if Cycle c.direction = Backward
5:      if Forward count ≥ Backward count then
6:          c.direction ← Forward
7:          Increment PC[e_i][e_j] for each forward direction pair from Cycle c
8:      else
9:          c.direction ← Backward
10:         Increment PC[e_i][e_j] for each backward direction pair from Cycle c
11:     end if
12: end for
```

There are $|V| = N$ redundant cyclic quorums for a given network. Each of these quorums has a corresponding cycle route in $Cycles$ causing Algorithm 1's For-loop on Line 2 to execute $O(N)$ times. Both forward and backward new pair counts, as well as, the final incrementing of PC after a direction is chosen requires enumerating all possible pairs in a particular cycle. From Equation 8, each quorum (i.e., cycle) will be of size approximately $O\left(\sqrt{RN}\right)$. Forming all pairs is $\binom{N}{2} = O\left(N^2\right)$ operation, so $O\left(\sqrt{RN}^2\right) = O(RN)$. All combined this results in an $O(N * RN) = O\left(RN^2\right)$ runtime to find the initial cycle direction.

The order of cycle iteration has an impact on the final direction of all cycles. It is possible that a cycle processed later may add pairs to PC that a previous cycle had already contributed. This opens up the possibility that a cycle processed earlier may be able to change to a more favorable direction and further reduce the number of missing pairs. Because of this, we have a second greedy heuristic, Algorithm 2.

Algorithm 2 Greedy Update Cycle Direction($Cycles$,V,PC)

```
1:  Changed ← True
2:  while Changed ≠ False do
3:      Changed ← False
4:      for all c ∈ Cycles do
5:          Decrement PC[e_i][e_j] for each pair from Cycle c
6:          Count number of new pairs added to PC if Cycle c.direction = Forward
7:          Count number of new pairs added to PC if Cycle c.direction = Backward
8:          if Forward count > Backward count then
9:              Increment PC[e_i][e_j] for each forward direction pair from Cycle c
10:             if c.direction ≠ Forward then
11:                 c.direction ← Forward
12:                 Changed ← True
13:             end if
14:         else if Backward count > Forward count then
15:             Increment PC[e_i][e_j] for each backward direction pair from Cycle c
16:             if c.direction ≠ Backward then
17:                 c.direction ← Backward
18:                 Changed ← True
19:             end if
20:         else
21:             Increment PC[e_i][e_j] for each pair from Cycle c
22:         end if
23:     end for
24: end while
```

This algorithm modifies the direction of a single cycle at a time to improve the missing pairs count. It continues modifying cycle directions until no further single cycle direction flips result in improvements being made. Although this appears to be an infinite loop on the surface, its runtime is considerably faster than evaluating all $O\left(2^N\right)$ possible combinations of cycle directions for a network.

Algorithm 2 uses a variable $Changed$ to determine if it should continue searching for a better combination of cycle directions (Lines 1-3). Line 4 iterates through all of the routed cycles. The pair count variable PC from the cycle direction initialization is utilized once again to account for the number of node pairs formed and also for how many pairs are missing. The cycle's direction is then chosen by whichever direction will eliminate more missing pairs from PC. Then all of that direction's pairs are added to PC (Line 9 or Line 15). If this selection causes the cycle's direction to change, then the $Changed$ variable is set to $True$ so the While-loop knows to continue searching. If the backward and forward missing pair counts are the same, then the direction defaults to the direction

The final publication is available at IEEE via http://dx.doi.org/10.1109/ICCCN.2016.7568595

the cycle is currently (Line 20). This process continues until an entire iteration of $Cycles$ results in no direction changes.

The outer While-loop continues until there are no direction changes. A direction change will only occur if it results in an improvement in the number of missing pairs. There are $|V| = N$ nodes and forming all pairs is $\binom{N}{2} = O\left(N^2\right)$ operation. So, even if every iteration of the While-loop only improved the number of missing pairs by 1, it would still only loop $O\left(N^2\right)$ times. The For-loop on Line 4 will execute $O\left(N\right)$ times as it iterates over all cycles. Both forward and backward new pair counts, as well as, the initial decrement and final incrementing of PC after a direction is chosen requires enumerating all possible pairs in a particular cycle, hence $O\left(RN\right)$, same as was calculated for Algorithm 1. All combined, this results in an $O\left(N^2 * N * RN\right) = O\left(RN^4\right)$ runtime to greedily update the cycle direction.

A 4th power is certainly not desired for most algorithms; but compared to $O\left(2^N\right)$, it scales well. For N greater than 16 this greedy approach is better than the brute force approach. And certainly for $N \leq 16$, the outer While-loop in Algorithm 2 will typically not behave in the worst case having to execute $O\left(N^2\right)$ times. In fact, given results from Section V-C1, on average less than 1% of the pairs were missing. Therefore, we can expect that the runtime behavior of the While-loop to be quite small for network inputs similar to those in our analysis.

B. Greedy missing pairs heuristic results

This section uses a similar experiment setup as Section V. We used the same four common networks (Fig. 2) and an implementation of the ECBRA heuristic [9] to perform the cycle routing. We also are using just a single cycle based on R redundant quorums sets.

1) Fault-free operational analysis: Section V-C1 showed significant resource usage reductions freeing up (wavelength) resource availability within the network and adding the potential for lower capital costs in terms of physical transmitters and receivers needed. The challenge, though, was this came at a cost. There were missing communication node pairs even in fault-free operation of the network.

Table II shows the significant improvements that our greedy heuristic had on this issue. An average reduction of greater than 94% of missing pairs for single cycle routing based on $R = 2$ redundant cyclic quorums. And a complete elimination of missing pairs for $R = 3$ redundancy.

Column 2 in Table II is the prior art paired cycle solution. This requires at minimum 38% more link resources on average, but no communication node pairs are missing. Originally in Section V, the cycle routing directly from the ECBRA heuristic was used. This produced the percent missing pairs.

The emphasis on the greedy heuristic was that chosen direction of the cycle matters. To emphasize this point and to confirm that our approach is actually doing something intelligent, we also included a random cycle direction algorithm (Column 4 and 8). There we see the random direction on average performs similar (and at times worse) than simply using the forward cycle directions directly from the ECBRA heuristic. Columns 5 and 6 (9 and 10 respectively) show the impact from our greedy heuristic. Nearly all of the missing pairs are removed on average from the $R = 2$ redundant quorum cycles and all of the missing pairs are eliminated from the $R = 3$ redundant quorum cycles. This makes the R redundant quorum, single cycle solution with greedy cycle direction heuristic a significant improvement over prior art's paired cycle approach. All communication pairs are formed in the network and done so with significantly fewer resources.

2) Fault-tolerant operational analysis: Section V-C2 illustrated a trade off between the significant reduction to resource usage and maintaining the same level of fault tolerance that prior art offered. Using only single cycles had a small impact on fault tolerance of the optical network, which for some networks and applications may still be acceptable given the improved resource usage. With the greedy heuristic addressing missing communication pair issues in the prior section, this section looks at the improvements to fault tolerance.

To model the fault(s), we simulate the failure of used edge(s), $(e_i, e_j) \in E$, in the 100 node mappings of each network model, $G = (V, E)$. The edges not used are ignored. We then examine the network's ability to serve all potential point-to-point requests. The results are then reported as fault coverage, total pairs able to communicate as a percentage of total point-to-point pairs. 100% would be perfect fault coverage, whereas 0% is no fault coverage at all.

Our single fault simulation results (Table III) shows the redundant quorum-based cycle technique with our greedy cycle direction heuristic had 97.92 - 99.67% and 98.59 - 99.84% fault coverages, $R = 2$ and $R = 3$ respectively, in the four networks tested. This was due to a 31.16 - 48.74% and 35.48 - 44.85%, $R = 2$ and $R = 3$ respectively, improvement over the number of missing pairs when only using forward cycle directions (i.e., Section V-C2). Again, we also test the performance against random cycle direction choices (Columns 4 and 7). This illustrates that the greedy heuristic, while it was designed to eliminate missing pairs for fault-free network operation, still provides an intelligent improvement for fault conditions as well.

With single link failures, the paired cycles had a mean missing communication pair rate of less than 3 pairs or less than 0.53% across all networks (95% CI). Hence the $R = 1$ (Paired) column shows mean fault coverage percentages is greater than 99.47% for all four networks. The redundant quorum cycles technique with our greedy heuristic, $R = 2$ and $R = 3$ (Single), reduced the number of missing pairs by greater than 30% when compared to without the heuristic (i.e., all forward cycles). Even with the improvement though, the fault tolerance still could not reach the level of coverage that the prior art's paired cycles. The Table III does show a competitive mean fault coverage rate (95% CI) of greater than 97.92 and 98.59%, respectively, across all networks when our greedy cycle direction heuristic was used with the single quorum cycle solution.

The final publication is available at IEEE via http://dx.doi.org/10.1109/ICCCN.2016.7568595

Table II
COMPARING FAULT-FREE OPERATION MEAN PERCENT MISSING NODE PAIRS (95% CI) BY SINGLE CYCLES USING REDUNDANT QUORUMS AND OUR GREEDY CYCLE DIRECTION HEURISTIC

Network	$R = 1$ (Paired) (%)	$R = 2$ (Single) (%)				$R = 3$ (Single) (%)			
		Forward	Random	Greedy	Reduction	Forward	Random	Greedy	Reduction
NSFNET	0.00 ± 0.00	0.95 ± 0.15	0.85 ± 0.12	0.02 ± 0.02	-98.27	0.04 ± 0.03	0.85 ± 0.12	0.00 ± 0.00	-100
ARPANET	0.00 ± 0.00	0.36 ± 0.07	0.31 ± 0.06	0.01 ± 0.01	-97.78	0.13 ± 0.04	0.31 ± 0.06	0.00 ± 0.00	-100
American	0.00 ± 0.00	0.49 ± 0.07	0.52 ± 0.07	0.02 ± 0.01	-95.56	0.21 ± 0.04	0.52 ± 0.07	0.00 ± 0.00	-100
Chinese	0.00 ± 0.00	0.27 ± 0.03	0.26 ± 0.02	0.01 ± 0.00	-94.98	0.09 ± 0.01	0.26 ± 0.02	0.00 ± 0.00	-100

Table III
COMPARING PERCENT MEAN FAULT COVERAGE (95% CI) OF THE SINGLE CYCLE, REDUNDANT QUORUM SOLUTION WITH OUR GREEDY CYCLE DIRECTION HEURISTIC EXPERIENCING A SINGLE LINK FAULT.

Network	$R = 1$ (Paired) (%)	$R = 2$ (Single) (%)			$R = 3$ (Single) (%)		
		Forward	Random	Greedy	Forward	Random	Greedy
NSFNET	99.47 ± 0.04	96.52 ± 0.09	96.84 ± 0.08	97.92 ± 0.07	97.81 ± 0.09	96.84 ± 0.08	98.59 ± 0.06
ARPANET	99.81 ± 0.01	98.25 ± 0.05	98.28 ± 0.05	98.79 ± 0.04	99.05 ± 0.04	98.28 ± 0.05	99.42 ± 0.03
American	99.60 ± 0.02	98.23 ± 0.04	98.21 ± 0.04	98.92 ± 0.03	98.90 ± 0.03	98.21 ± 0.04	99.34 ± 0.02
Chinese	99.90 ± 0.00	99.36 ± 0.01	99.38 ± 0.01	99.67 ± 0.01	99.71 ± 0.00	99.38 ± 0.01	99.84 ± 0.00

VII. CONCLUSION

The R redundant quorums significantly reduced resource usage and maintained high fault tolerance capabilities. The single cycle technique used significantly fewer resources (42.91 - 47.19% and 38.85 - 42.39% fewer, respectively), while at the same time maintained a high degree of fault tolerance with 96.52 - 99.36% and 97.81 - 99.71% fault coverage, respectively, on the single fault simulation.

With the paired cycles in prior art, the direction of communication pairs were not a consideration because both directions were always present. When we switched to single cycles, the resource usage fell dramatically (greater than 38%), and there was no longer the guarantee that both directions of a communication pair existed. To address this limitation, we developed a greedy algorithm to determine whether a cycle should be routed in the forward or backward direction in order to eliminate the most missing pairs and achieve higher fault tolerance. Missing pairs were reduced by over 94% for $R = 2$ redundant quorum single cycles and eliminated completely for $R = 3$. Fault tolerance improved as well. Missing pairs during single link failures decreased by over 30%, increasing fault tolerance to 97.92% - 99.67% and 98.59% - 99.84%, respectively.

ACKNOWLEDGMENT

Research funded in part by NSF Graduate Research Fellowship Program, IBM Ph.D. Fellowship Program, Symbi GK-12 and Trinect Fellowships at Iowa State University, and the Jerry R. Junkins Endowment at Iowa State University. The research reported in this paper is partially supported by the HPC@ISU equipment at Iowa State University, some of which has been purchased through funding provided by NSF under MRI grant number CNS 1229081 and CRI grant number 1205413. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the funding agencies.

REFERENCES

[1] D. Lastine, S. Sankaran, and A. K. Somani, "A fault-tolerant multipoint cycle routing algorithm (mcra)," in *Broadband Communications, Networks, and Systems*. Springer, 2012, pp. 341–360.

[2] D. Lastine, "Efficient communication using multiple cycles and multiple channels," Ph.D. dissertation, Iowa State University, 2014.

[3] C. J. Kleinheksel and A. K. Somani, "Scaling distributed all-pairs algorithms," in *Information Science and Applications (ICISA) 2016*. Springer, 2016, pp. 247–257.

[4] C. J. Kleinheksel and A. K. Somani, "Optical quorum cycles for efficient communication," *Photonic Network Communications*, pp. 1–10, 2015.

[5] C. J. Kleinheksel and A. K. Somani, "Enhancing fault tolerance capabilities in quorum-based cycle routing," in *Reliable Networks Design and Modeling (RNDM), 2015 7th International Workshop on*. Munich, Germany: IEEE, Oct. 2015, pp. 27–33.

[6] I. Chlamtac and A. Gumaste, "Light-trails: A solution to ip centric communication in the optical domain," in *Quality of Service in Multiservice IP Networks*. Springer, 2003, pp. 634–644.

[7] J. Fang, W. He, and A. K. Somani, "Optimal light trail design in wdm optical networks," in *Communications, 2004 IEEE International Conference on*, vol. 3. IEEE, 2004, pp. 1699–1703.

[8] Y. Li, J. Wang, A. Gumaste, Y. Xu, and Y. Xu, "Multicast routing in light-trail wdm networks," in *Global Telecommunications Conference, 2008. IEEE GLOBECOM 2008. IEEE*. IEEE, 2008, pp. 1–5.

[9] A. Somani, D. Lastine, and S. Sankaran, "Finding complex cycles through a set of nodes," in *Global Telecommunications Conference (GLOBECOM 2011), 2011 IEEE*, Dec 2011, pp. 1–5.

[10] G. P. Agrawal, *Nonlinear fiber optics*. Academic press, 2007.

[11] W. D. Grover and G. Shen, "Extending the p-cycle concept to path-segment protection," in *Communications, 2003. ICC'03. IEEE International Conference on*, vol. 2. IEEE, 2003, pp. 1314–1319.

[12] A. Khalil, A. Hadjiantonis, G. Ellinas, and M. Ali, "Pre-planned multicast protection approaches in wdm mesh networks," in *Optical Communication, 2005. ECOC 2005. 31st European Conference on*. IET, 2005, pp. 25–26.

[13] W.-S. Luk and T.-T. Wong, "Two new quorum based algorithms for distributed mutual exclusion," in *Distributed Computing Systems, 1997., Proceedings of the 17th International Conference on*. IEEE, 1997, pp. 100–106.

[14] L. Tang, W. Huang, M. Razo, A. Sivasankaran, M. Tacca, and A. Fumagalli, "Multicast tree computation in networks with multicast incapable nodes," in *High Performance Switching and Routing (HPSR), 2011 IEEE 12th International Conference on*. IEEE, 2011, pp. 95–100.

The final publication is available at IEEE via http://dx.doi.org/10.1109/ICCCN.2016.7568595